Sonnox

Oxford SuprEsser

User Guide

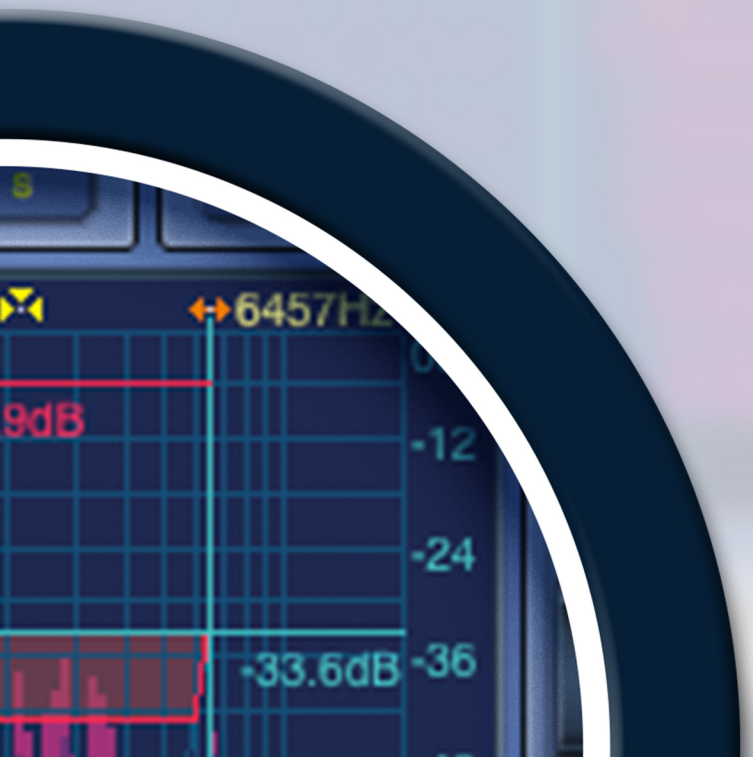

Contents

1 Introduction 3

2 Main Features 5

3 Operation 6

3.1 Signal Processing Overview 6

3.2 Resolution, Kernel Sizes, and Delays 8

3.2.1 Delay Compensation and Audio Buffer Sizes 10

3.2.2 Over-Taxing the Host 11

3.3 Advanced Operation 12

3.3.1 Trigger and Gain-Ducking Modes 12

3.3.2 Automatic Level Tracking 15

3.3.3 Bandpass Filter Modes 16

4 Description of Controls 18

4.1 Basic Screen Controls 18

4.1.1 Touch Pad Controls 18

4.1.2 Options Menu 20

4.1.3 Input Monitor Section 21

4.1.4 Listen Section 21

4.1.5 FILTERS Section 22

4.1.6 DYNAMICS Section 23

4.1.7 OUTPUT MONITOR Section 24

4.2 Advanced Screen Controls 25

4.3 Graphical Screen Controls 29

4.3.1 Explanation of Controls 30

5 Oxford SuprEsser DS (AAX DSP) 32

6 Specifications 33

6.1 SuprEsser DS – Pro Tools | HDX & S3L – Instances per chip 34

7 Preset Manager Toolbar 34

8 Supported Platforms 35

9 System Requirements 35

10 Copyright and Acknowledgements 36

1 Introduction

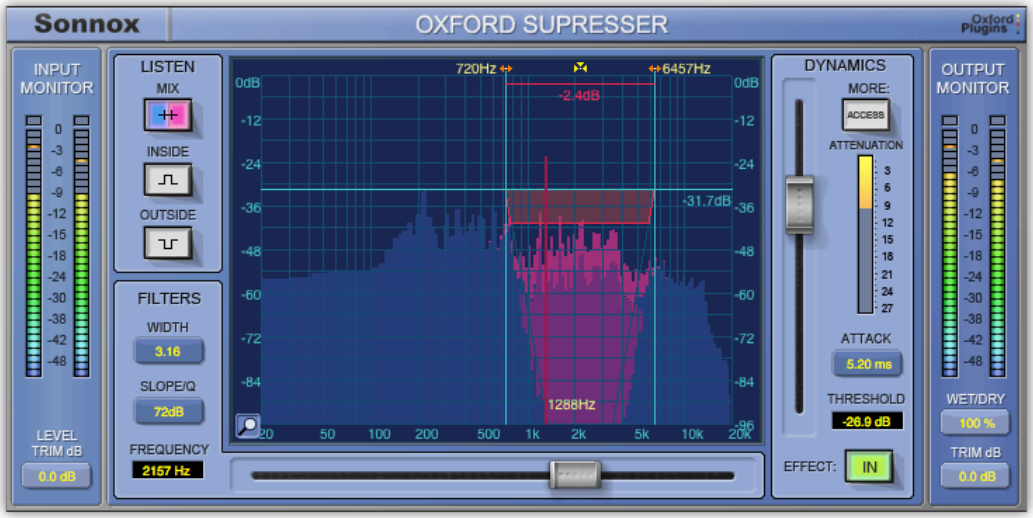

The Sonnox Oxford SuprEsser was designed primarily to be the last word in ‘de-essing’ applications. While we worked on creating the best possible de-essed sound, we found that we needed more control than was available on conventional De-Essers. Most de-essing work can be carried out in ‘simple’ mode but sometimes, to do the job properly, it is necessary to utilise all the controls of the underlying engine — a full-blown dynamic EQ, or frequency- conscious compressor.

Thus a new concept was born — a simple and intuitive De-Esser, plus access to a much more sophisticated frequency-specific compressor, giving complete control over aggressive frequencies wherever they may be in the audio spectrum. With the Oxford SuprEsser you can take out unwanted frequency peaks as and when they occur — unlike applying a fixed EQ to the whole channel.

As a De-Esser, the Oxford SuprEsser is designed for the treatment of sibilance and fricatives in vocals, and the treatment of unwanted whistles and ’spirant’ artefacts associated with wind instruments. However, it can equally be applied to removing low-end plosives and thuds from over-close vocal work, without affecting nearby components in the frequency spectrum, keeping the low-end intact. For the most natural sounding results, the Oxford SuprEsser de- esses only the frequency band you select — so you won’t end up with an over-de-essed lisp-like voice with all the high frequencies gone!

Detailed visual feedback is provided by a highly intuitive graphical display, allowing quick identification of the frequencies that need treatment, and where to set the threshold. The threshold level and peak-hold level of the user-definable

band are shown on the graph, alongside the FFT (Fast Fourier Transform) display of the narrow band signal, which includes retention of the peak level and the frequency containing the most energy.

At the heart of the Oxford SuprEsser is an enhanced version of the compressor section of the Sonnox Oxford Dynamics plug-in, which is renowned amongst professional users for its consistent delivery of the precise and transparent control of peak signals. Around this is built a pair of crossover filters to make the compressor react only to the defined frequency band. These linear-phase filters are modelled on the filters from the Sonnox Oxford EQ, making the Oxford SuprEsser useful for precise mastering as well as mixing work.

Three listen modes allow the user to listen to the **Mix**, the output of the bandpass filter (**Inside**), or the output of the band-reject filter (**Outside**).

Careful thought has been put into making the Oxford SuprEsser extremely easy and quick to use. The screenshot on the title page shows the plug-in as it appears when first activated. Once the frequency band has been defined, simply lower the threshold fader until the gain reduction meter starts to kick in. The plug-in then automatically tracks the general signal level and the threshold follows accordingly, so that it gives the same relative amount of gain reduction as the signal level rises or falls. This is perfect for vocals where, for example, a vocalist is louder in the chorus than the verse, but you want to apply the same relative amount of de-essing, but don't want to over de-ess in the chorus. This Auto-Level-Tracking mode lets the Oxford SuprEsser do all the hard work. It can also be switched off for a fixed threshold.

In its default mode of operation, the Oxford SuprEsser feeds the defined **Band** signal to the compressor to affect only this narrow-band signal. The result is that the compressor reacts only to specific frequency components when they reach a specific threshold, and applies compression only to these specific frequency components, leaving the rest of the spectrum untouched. Audio **Wide** mode can also be selected to allow full-band compression reacting to only specific frequencies, or any combination of **Band** and **Wide** defining when the compressor reacts, and what it compresses.

The Oxford SuprEsser comes complete with many presets to act as good starting points; alternatively, the advanced section provides full access to all controls for precise correction, or even creative use.

2 Main Features

- Highly featured professional De-Esser
- Linear-phase Dynamic EQ
- Transparently controls aggressive frequencies
- Automatic level-tracking follows energy level (eliminating the need to automate threshold)
- Large intuitive graphic display makes finding frequencies very easy
- Full spectrum operation (20Hz–20kHz)
- Three different ‘Listen’ modes
- Very easy to use
- Advanced mode for ultimate control of the Dynamic EQ
- Many creative as well as corrective uses
- Presets provide good starting points

3 Operation

3.1 Signal Processing Overview

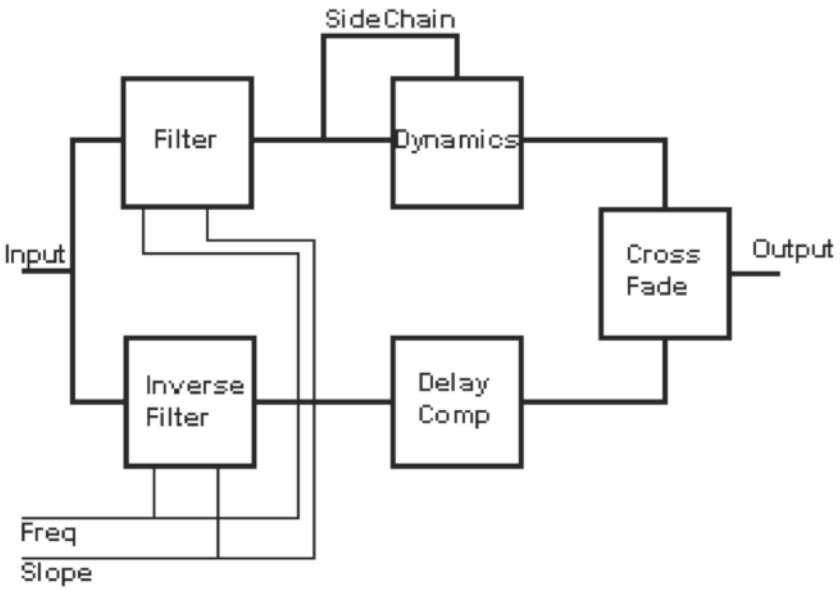

Basic SuprEsser signal flow

The Oxford SuprEsser contains a pair of mutually opposing filters — by default one is a narrow bandpass filter, the other is the complementary narrow band-reject filter. This results in one signal path containing just the contents of the band of interest, and another signal path containing the input signal with this band entirely removed; when mixed back together in equal ratios, you get the original signal.

The bandpass signal is usually fed to the compressor, both to the sidechain and to the main signal path, so that it is just this signal that triggers gain reduction, and it is just this signal that is affected by any gain reduction.

The bandpass/reject filters will switch automatically to other EQ types when conditions warrant it, giving a total of four EQ types. For example, when the 'Width' control narrows the bandpass filter to its minimum, a High-Q notch filter is invoked that provides a much better reduction of a very narrow band of frequencies. When either side of the band window touches the end stops (20Hz or 20kHz), the bandpass filters change to LF-Cut or HF-Cut as necessary.

Advanced Mode Signal Flow

In addition to the four EQ types, there are four different compressor modes that are concerned with which signal is passed to the sidechain and which signal is passed to the main input of the compressor. See Section 2.3 Advanced Operational Modes for more information on this.

The following diagram provides a more complete signal flow to illustrate the elements required for advanced modes of operation:

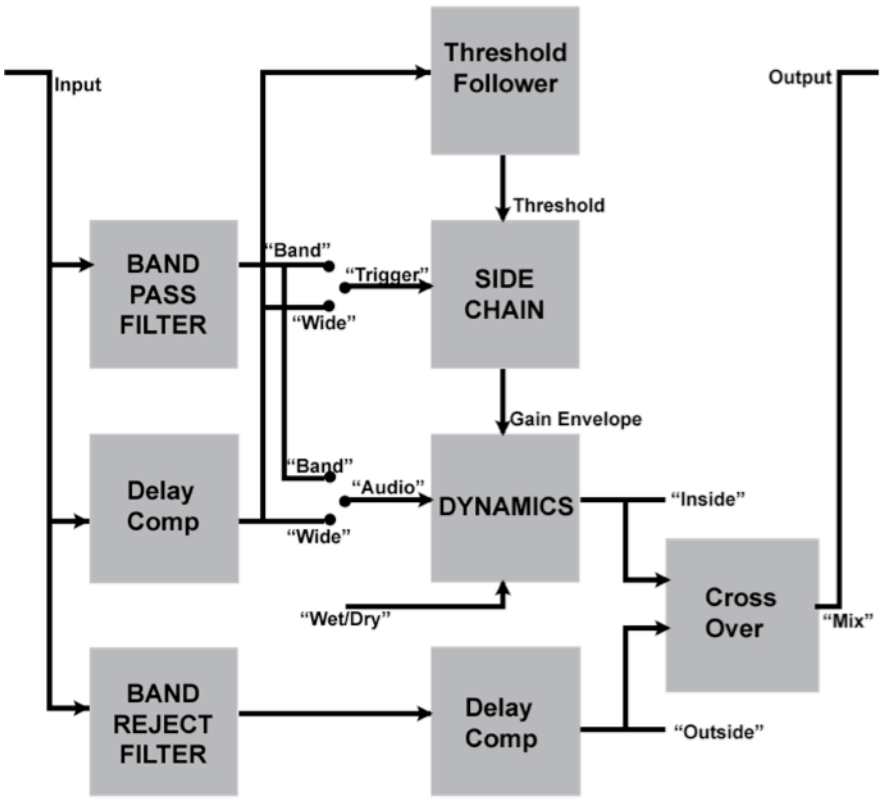

Advanced SuprEsser signal flow

The 'Wet/Dry' control is a frequently requested feature that allows you to add back the uncompressed signal to a highly compressed signal to add in some punch. As you can see from the diagram above, the Wet/Dry blend control is implemented inside the dynamics module, and therefore only operates as expected when you are listening to the 'Mix', or 'Inside'.

3.2 Resolution, Kernel Sizes, and Delays

The linear-phase filters used by the Oxford SuprEsser require an 'Impulse Response Kernel' to model the response of the internal Oxford Filters. The size of this kernel (as measured in samples) determines both the plug-in delay and the accuracy of the model, which in turn affects the performance, especially at lower frequencies.

Large kernels

Better performance at low frequencies, but longer overall plug-in delays

Small kernels

Adequate performance at high frequencies, and shorter plug-in delays

While small kernels are adequate for high frequency work such as de-essing, the performance at low frequencies will cause poorly defined filter slopes, and poor separation of 'Inside' from 'Outside'.

We have found that host applications are generally not able to readily adjust their delay compensation engines if the kernel size is adjusted dynamically at run time. The Native Oxford SuprEsser is therefore released as three separate plug-ins; each with a different fixed kernel size.

Oxford SuprEsser

This is the standard version, and has a kernel size of 2048 samples. This is suitable for work across the entire frequency range at 44.1kHz sample rate, and thus is the standard plug-in to reach for, especially for the mixing stage of a project. Depending on your audio buffer size, the delay or latency will be somewhere in the region of 1044 to 3092 samples. See the section below for information on reducing the delay.

Oxford SuprEsser HR/HighRes

This is the large kernel version, with a kernel size of 8192 samples. This version gives superior resolution at the lowest frequencies, and thus is particularly suitable for low frequency mastering work. Remembering that, as you increase the sample rate, the resolution at the low end will be correspondingly reduced, this version is particularly suitable for use at higher sample rates, especially 176.4kHz or 192kHz. This version will have a very significant delay that can be beyond the ability of Pro Tools HD to automatically compensate for. The true delay is reported correctly in the track delay information, and depending on your audio

buffer size, the delay will be somewhere in the region of 4116 to 12308 samples. See the section below for information on reducing the delay.

Oxford SuprEsser LL/LowLatency

This small kernel version has a kernel size of 512 samples.This version gives superior performance in terms of having a small delay, and thus is more suitable for live de-essing work at low sample rates, and for the tracking phase of a project, where you are laying down new tracks from midi instruments, and you don't want a significant delay between what you are playing and what you are hearing. The smaller kernel size means the resolution at low frequencies will be poor, and so cannot be used much below 400Hz. Depending on your audio buffer size, the delay or latency will be somewhere in the region of 276 to 3072 samples. See the section below for information on reducing the delay.

Oxford SuprEsser DS (AAX DSP)

This version has an even smaller kernel size of 128 samples, in order to further reduce the processing delay. Due to this, the filter range has been reduced to 1 kHz minimum, in order to prevent poor low frequency performance from producing unexpected results.

For further information, see Section 5.

3.2.1 Delay Compensation and Audio Buffer Sizes

The plug-in delay produced by the Oxford SuprEsser depends both on the kernel size/resolution and the audio block size.

The block size is the size of the sample buffers passed to the plug-in by the host, and is usually specified in your audio hardware preferences/configuration/setup page. The reason that the plug-in delay depends on the block size is that the plug-in must accumulate a whole kernel sized block of samples before it can process them.

In order to ensure the minimum plug-in delay, make sure the block/buffer size is the same as or greater than the kernel size/resolution setting!

For example, if the kernel size is set to 512, ie. you are using the Low Latency version, and if your block/buffer size is 512 or 1024, this will ensure that the plug-in produces the minimum delay of 276 samples.

If you use a smaller block size than the kernel size, the overall delay of the plug-in will go up, not down. For example, if the block size is 256, then the Low Latency version will produce a delay of 532 samples.

3.2.2 Over-Taxing the Host

The Oxford SuprEsser uses a process called convolution to implement filtering, a process that is expensive on CPU resources. When using very large kernel sizes, in combination with small audio buffer sizes, it is possible for the plug-in to take longer than an entire audio block to complete processing, with the result that (from a monitoring point of view) the playback breaks up. This can manifest as obvious clicks and pops, or more subtly as occasional quiet clicks.

If you experience clicks, keep in mind that these are monitoring only. They will not be present when bouncing down your mix.

If you experience clicks, increasing your audio buffer size will resolve the issue. Once again, the best buffer size for the Oxford SuprEsser is the same as the kernel size (or resolution). Generally speaking, the only reason for small audio buffer sizes is for tracking, or live working requirements.

The plug-in issues a red warning label next to the Sonnox button when you are using a combination of small buffer size and large kernel size that we have found to generally result in clicks. This is not a reliable indicator, however; it completely depends on your system CPU speed, sample rate, and other factors.

3.3 Advanced Operation

3.3.1 Trigger and Gain-Ducking Modes

There are four major operational modes selectable via the TRIGGER and AUDIO buttons in the Advanced Mode (MORE:) section, as described below:

Mode:		Signal to:		Result
Trigger	Audio	Comp	Side Chain	
Band	Band	Filtered signal	Filtered signal	Gain ducking occurs only in the narrow band, triggered by the narrow band.
Band	Wide	Delayed input signal	Filtered signal	Gain ducking of (wide) input signal, triggered by narrow band signal. Equivalent to compressor in side chain EQ mode.
Wide	Band	Filtered signal	Delayed input signal	Gain ducking occurs only in the narrow band of the spectrum, triggered by the wide input signal.
Wide	Wide	Delayed input signal	Delayed input signal	Equivalent to compressor with no side chain EQ.

Band-Band Mode

In this mode (the default) both the trigger signal (to the compressor sidechain) and the main compressor signal is the narrow-bandpass-filtered signal.

The output of the compressor is then fed to the crossover block (the 'Listen' section) for mixing back in with the version of the input signal that has been filtered with the corresponding narrow-band-reject filter. The result is that the plug-in affects only a narrow frequency band, triggered by that same frequency band, and does not affect the rest of the frequency spectrum.

In this mode, when you also use the Auto Level Track function (AUTO IN), the threshold follows the general level of the signal post the band-reject filter, (ie. everything except the troublesome frequencies).

Band-Wide Mode

In this mode, the trigger signal (to the compressor sidechain) is the narrow-bandpass-filtered signal, and the main compressor signal is the (wide) delayed input signal.

This means that when the narrow band signal triggers gain ducking, the ducking occurs over the whole frequency spectrum, as in traditional sidechain EQ compression.

This mode offers one major advantage over the equivalent mode in the Oxford Dynamics compressor — you have a WET/DRY control on the output stage, so you can mix the original signal back into the compressed signal.

In this mode, when you also use the Auto Level Track function (AUTO IN), the

threshold follows the general level of the signal post the narrow-band-reject-filter, (ie. everything except the troublesome frequencies).

Furthermore, the crossover controls (in the LISTEN section) have no effect on the signal because the output of the compressor is already a complete wide-spectrum signal.

Wide-Band Mode

In this mode, the trigger signal (to the compressor sidechain) is the delayed input signal, and the main compressor signal is the narrow-bandpass-filtered signal.

This mode is useful, for example, when a broad-spectrum trigger is required to trigger reduction of a specific frequency. One example could be to process a kick drum/bass combination to reduce a rattle or click whenever the kick occurs.

In this mode, when you also use the Auto Level Track function (AUTO IN), the threshold follows the general level of the delayed input signal.

Wide-Wide Mode

In this mode, the trigger signal (to the compressor sidechain) and the main compressor signal are both fed from the delayed input signal. In other words, you now have an ordinary compressor with no frequency specific sidechain processing.

This mode offers one major advantage over the Oxford Dynamics compressor — you have a WET/DRY control on the output stage so that you can mix the original signal back into the compressed signal. This can be used to create that unique

sound in which a fully compressed signal (with no dynamic headroom) has some punch added back into it.

In this mode, the LISTEN section has no effect on the signal because the output of the compressor is already a complete wide-spectrum signal.

3.3.2 Automatic Level Tracking

Auto Level Tracking mode is enabled by selecting the LEVEL TRACKING/AUTO IN button. This is ON by default.

The purpose of this mode is to automatically adjust the threshold level to follow the general signal level of the wide-band input. This means that if some vocals wander from loud to quiet, the same amount of relative gain reduction is applied when a transient peak occurs above the general level.

Generally speaking, once you have isolated as narrowly as possible a troublesome sound in the frequency spectrum, you will see the red peak-hold marker, inside the bandpass filter on the graphic display, indicating the peak level of the troublesome sound. Now enable Auto Track mode, and bring the THRESHOLD level down until its corresponding line on the graph is below the red line. You will then see the gain reduction meter (ATTENUATION) starting to show gain reduction.

If the general signal level of the material changes, the threshold will now follow that level, so that the gain reduction meter should continue to indicate the same amount of reduction of the peaks.

The algorithm to track the general signal level is not as simple as it appears. First, one of the most important characteristics of speech is its staccato nature, and the frequent pauses or silence. If the threshold truly followed the signal level, each time there was a pause broken by a word starting with a consonant, the threshold would have fallen too low, and the consonant would be over gain-reduced before the threshold recovered its normal level.

To get around this challenge, the algorithm implements a working range of 24dB, and any sample value outside this window is not included in the calculation of the general signal level. If you were to intermittently mute the input signal, you will see that the threshold effectively just stays where it was when there was an active signal, and continues tracking again when the active signal is restored. Yet if you slowly fade out the input signal level, the threshold will follow, only giving up after 36dB of gain reduction.

Auto Level Tracking therefore works well with continuous material, or with conversation broken up with frequent pauses.

Setting the After-Silence Start Level

If the above algorithm is tending to under-correct when the vocalist starts singing after a pause, you can, if you wish, set the 'After-Silence Start Level' by moving the threshold fader while the plug-in is receiving silence or is non-active.

If you move the threshold level while there is silence or no signal, this will tell the plug-in what level to assume when the signal comes back after going away for a while. This level will remain indicated on the graph as a ghost line.

3.3.3 Bandpass Filter Modes

Bandpass Mode

By default, the Oxford SuprEsser uses a bandpass filter (along with its inverse filter, the band-reject filter) to isolate the audio in the frequency range selected. The FREQUENCY and WIDTH controls defines the low and high edges of the filters used, and the SLOPE/Q control defines how quickly a signal blends from gain-ducked to non-gain-ducked in the frequency spectrum. The effect of less steep slopes is to give a smoother blend between gain- reduced and non-gain-reduced regions of the frequency spectrum, at the cost of a less specific trigger.

In this mode, the bandpass filter is created by using pairs of Oxford R3 Filters, giving a total of up to 72 dB per octave of separation.

There are three other modes possible, which automatically switch into operation in certain situations:

High-Q Notch Mode

When you reduce the WIDTH control all the way down to 0.2 octaves, you activate High-Q Notch mode. In this mode, instead of using R3 Filters, the plug-in uses four Oxford EQ Type-2 Filters, giving a total of 80dB of gain reduction at the centre point of the notch. This mode is useful when the band of energy you are interested in reducing is extremely narrow, or virtually a single frequency, like a whistle.

In this mode, if desired, you can reduce the Q of the filters by reducing the SLOPE control. The effect of this is to give a smoother blend between gain-reduced and non-gain-reduced regions of the frequency spectrum, at the cost of a less specific trigger.

LF-Cut Filter Mode

When you reduce the left edge of the bandpass filter all the way down to 20Hz, you activate LF-Cut Filter mode. This means that the lower edge of the bandpass window is effectively at 0Hz, and this is useful, for example, when working on eliminating sub-bass plosives and thuds/booms that contain DC components.

In this mode you can think of operation as an LF-Cut Filter that activates only when the signal reaches a certain threshold, in which the cut-off frequency is defined by the upper edge of the bandpass window.

HF-Cut Filter Mode

When you increase the right edge of the bandpass filter all the way up to 20 kHz, you activate HF-Cut Filter mode. This means that the upper edge of the bandpass window is effectively at infinity Hz, and this is useful, for example, when you want to duck the entire HF part of the signal.

In this mode you can think of operation as an HF-Cut Filter that activates only when the signal reaches a certain threshold, in which the cut-off frequency is defined by the lower edge of the bandpass window. This is how some more primitive De-Essers work.

4 Description of Controls

4.1 Basic Screen Controls

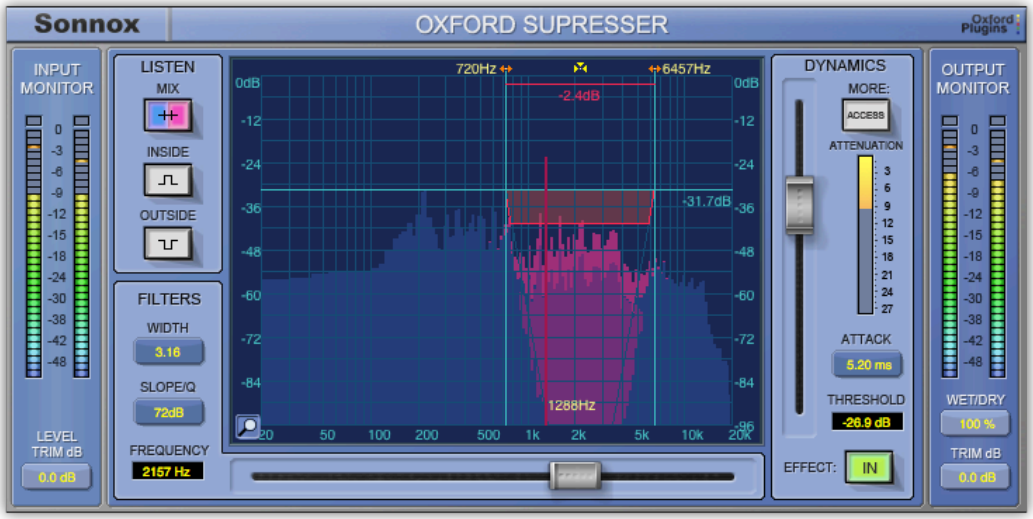

When you insert the plugin, the display defaults to a zoomed out mode, so you can see immediately the overall picture of where the centre band is in the spectrum.

Whenever you move the centre band in any way, the centre frequency of the band is clearly displayed (in yellow) at both the top and bottom of the display in an easy-to-read format. This is so that those who are aiming for a particular frequency can see the centre frequency clearly as they move towards it. Experienced engineers often recognise the frequency they wish to subdue, and so want clear guidance to aim for that frequency.

4.1.1 Touch Pad Controls

Many of the controls on the Oxford SuprEsser are implemented using ‘Touch Pads’. These give a clear display of the SI value of a control, and respond to the following actions:

- Left-Click and drag up/down: Increase/decrease the value
- Right-Click and drag up/down: Fine adjustment of value
- Double Click: Directly edit value
- Scroll wheel up/down: Increase/decrease value

In VST and Audio Unit hosts

- Shift + Left-Click and drag up/down: Fine adjustment of value
- Control/Command + Left-Click: Set to default
- Shift + Scroll wheel: Fine adjustment of value

In Pro Tools

- Command + Left-Click and drag up/down: Fine adjustment of value
- Alt + Left-Click: Set to default
- Command + Scroll wheel: Fine adjustment of value

Other controls (for example, faders, buttons and the graphic display) are generally consistent with the rules outlined above for touch pad control.

Most controls have an explanatory fly-out help window that is activated by hovering your mouse over the control for two seconds.

4.1.2 Options Menu

Clicking the Sonnox button produces a drop-down options menu:

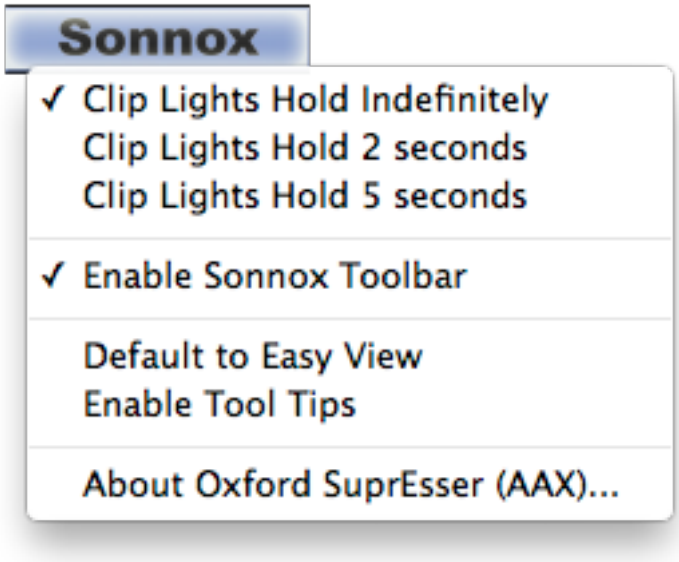

Clip Lights

These options determine the approximate time that an overload indicator will stay on for when the plug-in has detected a full-level sample at either its input or output.

Enable Sonnox Toolbar

Displays or hides the Sonnox Preset Manager Toolbar (see Section 4).

Show Preset Name Path

Shows the hierarchical directory path for Presets that are stored in sub-folders of the default Preset folder.

Default to Easy View

Determines whether the plug-in will default to ‘Easy’ or ‘Advanced’ mode on opening.

Enable Tooltips

Enables/disables the fly-out help windows that are activated by hovering the mouse over a control for two seconds.

Abbout Sonnox Oxford SuprEsser

Displays the date, version and build number of the plug-in.

4.1.3 Input Monitor Section

Input Level Meter

This meter is designed to give exactly 1dB per LED for the top 18dB of dynamic range, and 2dB per LED thereafter. This gives a clear and intuitive impression of the working headroom.

Level Trim dB Touch Pad

This allows you to adjust the input signal level by up to +/- 12dB.

Although it is true that, in default operation, the threshold of the compressor is more or less level independent, you may wish to adjust the input level to bring peaks down below 0dB, because the threshold cannot go above 0dB.

4.1.4 Listen Section

This section of the plug-in allows you to control what you are listening to in terms of the crossover that mixes the outputs of the two signal streams.

Some users prefer to isolate the troublesome sound, and keep narrowing the bandpass filter to just contain it — this way the compressor affects only the absolute minimum of the frequency spectrum. You can do this by clicking on the INSIDE button, labelled with a bandpass symbol.

Others prefer to work in band-reject mode, where you listen to the narrow band-reject-filtered signal and narrow the band until you just begin to hear the troublesome sound. You can do this by clicking on the OUTSIDE button, labelled with the band-reject symbol.

MIX Button

This sets the cross fade exactly to 50%, providing an equal mix of band-reject signal and (compressed) bandpass signal, which will be sent to the output.

INSIDE Button

This sets the crossfade so that you are listening to output of the bandpass filter, as processed by the compressor. You can then use the EFFECT: IN button to listen with or without compression, or use the WET/DRY control to adjust the blend between compressed and non-compressed signal.

In this mode, if you sweep the centre frequency of the band up and down the frequency spectrum, you hear only a narrow part of the spectrum, and so it is easy to home in quickly to locate exactly where in the spectrum a troublesome noise is.

OUTSIDE Button

This sets the crossfade so that you are listening only to the output of the band-reject filter — which is everything outside of the bandpass window, hence the name. In this mode, if you sweep the centre frequency of the band filter up and down the spectrum, the troublesome noises will disappear when you have hit the area to work on.

4.1.5 FILTERS Section

This section allows you to speficy the filter characteristics:

FREQUENCY fader and touch Pad

The horizontal fader below the graphical display, and its associated touch pad in the Filters section, control and display the centre frequency of the bandpass and band-reject filter pair. The control range goes all the way down to 20Hz to allow for low frequency treatment as well as mid and high frequency treatment. The centre frequency of the band filters can also be changed on the graph itself by dragging (horizontally) the yellow 'centre' symbol at the top of the display. See Section 3.3 Graphical Screen Controls if you are unclear about the location of this latter control.

When the low edge of the bandpass filter reaches 20Hz, the plug-in switches over to using just the high-edge crossover, which means the bandpass window extends all the way down to DC, providing a new useful mode. See 2.3 Advanced Operational Modes for more information.

WIDTH Touch Pad

This controls, in octaves, the width of the bandpass and band-reject filter pair. The control range is wide enough to allow almost the entire spectrum from 20Hz to 20kHz. You can also change the filter width, also relative to its centre point, with vertical drags of the graph's yellow centre, or by dragging horizontally on the filter's upper or lower frequency markers (see Section 3.3 Graphical Screen Controls for more details).

When you reduce the width down to 0.2 octaves, this triggers the bandpass filter to switch to a new type of filter — a High-Q notch filter. This is useful when a very narrow frequency band needs removal, such as a whistle at a particular frequency. See Section 2.3 Advanced Operational Modes for more information.

SLOPE/Q Touch Pad

This controls the slope of the filters used to implement the bandpass and

band-reject filters. Generally speaking, you want the highest setting for maximum separation between troublesome and non-troublesome sounds. However, you may find musical or creative possibilities in using gentler slopes. For example, when using lower settings there is a smoother blend between gain-ducked components and the original, which will be more noticeable when using large amounts of gain reduction.

The slope can also be changed on the graph itself by dragging vertically on the red arrows at the top of the filter's left and right width lines (see Section 3.3 Graphical Screen Controls for more details)

4.1.6 DYNAMICS Section

THRESHOLD Fader and Touch Pad

The THRESHOLD Fader is linked to a cyan threshold line on the graph, where it has a relationship with the red peak level line. The THRESHOLD may also be changed by dragging vertically on its touch pad.

Generally speaking, after you have isolated as narrowly as possible a troublesome sound within the frequency spectrum, using the bandpass filter adjustment controls (Frequency, Width, or the drag controls on the graph), you should then lower the threshold fader until the gain reduction begins to kick in.

By default, the threshold is relative to the general signal level, and this is why the threshold line moves up and down on the graph as the level of the audio signal changes. You should aim to move the threshold until you hear the right amount of gain-reduction taking place. The plug-in, by tracking the general signal level, attempts to keep this gain reduction amount constant as the signal level climbs or falls.

EFFECT: IN Button

This enables/disables dynamics processing of the main compressor signal, and allows you to do glitch-free with/without comparisons of the gain-reduction effect.

ATTACK Touch Pad

The Attack control quickens or slows down the reaction time of the compressor. Smaller numbers mean faster reaction times and harder sounding results. Go too fast and you may hear distortion when working at lower frequencies. Bigger numbers mean slower reaction times and softer sounding results. Go too slow and you may not react to fricatives quickly

enough to catch the hard front edges of consonants like T's.

ATTENUATION Meter

This meter gives an indication of how much gain reduction is occurring, both instantaneously, and with a peak-hold level that indicates the maximum reduction in the previous couple of seconds. The meter operates in increments of exactly one dB.

MORE: ACCESS Button

When you need to access the complete set of dynamics controls, click on the ACCESS button to open up the 'advanced' control set. This includes many more controls than would generally be needed for normal operation, controls that would otherwise clutter the primary functionality of the plug-in. Nonetheless, these controls provide the advanced or curious user with many more creative, or tweaking, possibilities.

The ACCESS button is to be used especially if you want to change to an absolute rather than a relative threshold, or if the time constant of the threshold signal tracking is too fast or too slow. See Section 2.3 Advanced Operation, for more information.

4.1.7 OUTPUT MONITOR Section

TRIM dB Touch Pad

This allows you to reduce the output level by up to 12dB. Dithering is applied after the output gain control, so it may be necessary to reduce this value by a small amount to avoid clipping.

WET/DRY Touch Pad

This control enables you, if you want, to mix in some of the original input with the output, rather like the Wet/Dry control on a reverb. The use of this is to allow you to generate a very fat compressed sound, and then add a little of the original back in to regain some dynamic character. This is particularly useful when operating the Oxford SuprEsser as a basic compressor (ie. in 'Wide/Wide' mode — see Section 2.3 Advanced Operation, for more information).

At 100%, you hear only the post-dynamics signal; at 0% you hear only the input. At 50%, you hear an equal blend of the two. You will normally want this set to 100%.

Output Level Meter

The output meter is designed to give exactly 1dB per segment for the top

18dB of dynamic range, and 2dB per segment thereafter. This gives a clear and intuitive impression of the working headroom and dynamic range. There is a peak-hold feature that holds the highest peak in the last two seconds, helping to give you a better impression of the working dynamic range — the number of segments between the peak and general level.

4.2 Advanced Screen Controls

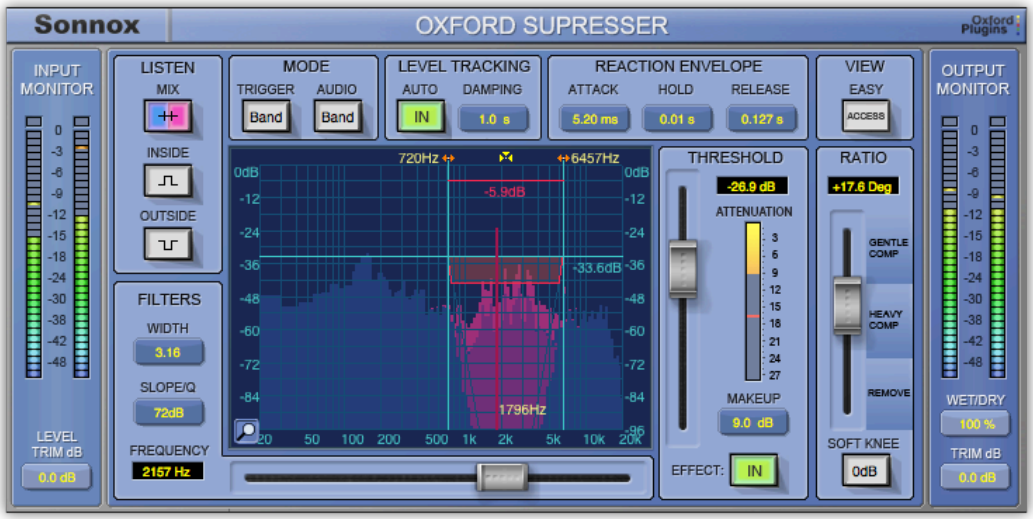

The advanced screen, and its associated controls, is provided to allow the Oxford SuprEsser to be used as a general-purpose frequency-specific compressor, giving you access to all the controls within the plug-in's dynamics section.

MODE Section

TRIGGER Button

This toggles between Wide and Band, and controls the type of signal fed to the sidechain of the compressor, ie. the signal that will be used to trigger the gain ducking. See Section 2.3 Advanced Operational Modes for a description of the modes.

AUDIO Button

This button also toggles between Wide and Band, and controls the type of signal fed to the main signal path of the compressor, ie. the signal that will be gain-ducked. See Section 2.3 Advanced Operational Modes for a description of the modes.

LEVEL TRACKING Section

AUTO IN Button

This controls the mode of the threshold functionality. In 'Auto-Level-Tracking' mode (ie. with the IN button selected), the compressor's trigger threshold follows the general level of the wide input signal, and the THRESHOLD fader adjusts the threshold relative to that general signal level. If you want the same relative level of gain reduction for a quiet sound as for a loud sound (for example, in each case you want 6dB of gain reduction relative to the general signal level), then use Auto-Level-Tracking mode, and set the THRESHOLD fader 6dB lower than the peaks shown on the graph.

In 'Absolute' mode (ie. with the IN button not selected), the THRESHOLD fader controls the absolute level of the compressor's trigger threshold — so if it is set to -10dB, signals over that level will trigger gain reduction.

DAMPING Touch Pad

In 'Auto-Level-Tracking' mode, this controls the time constant used for following the general level of the wide input signal. In other words, the smaller the damping value is, the more quickly the threshold level will rise or fall to a new level; the larger the value, the more stable the threshold value.

Generally speaking, you want to have the largest value possible consistent with the material you are working with. For general dialogue or vocals, a larger value will be adequate because there is not so much variation in general signal level. For a speaker who rapidly changes from mumbling to shouting, then a smaller constant will be required for the compressor to adapt quickly enough.

The algorithm used to implement 'Auto-Level-Tracking' responds more quickly to rises in the input level, and more slowly to drops in level. This means that the plug-in will tend to under-gain-reduce when the level drops rapidly, and minimise over-gain-reduction when the level rises rapidly. This helps to compensate for the nature of vocals and speech in which the troublesome sounds often leap out of silence.

REACTION ENVELOPE Section

The ATTACK, HOLD and RELEASE controls together define the shape of the gain reduction envelope created when a transient breaches the threshold.

ATTACK Touch Pad

This defines how quickly the gain reduction can kick in. While fast attack times often seem to be the best idea, slower attack times can sound more musical and natural. In general, for de-essing, fast attack times are required — as fast as possible without distortion. When working with low frequency sounds, however, you may want slower attack times.

HOLD Touch Pad

This defines how long the gain reduction holds its level before it is allowed to decay away. Generally speaking this needs to be as small as possible for de-essing work.

RELEASE Touch Pad

This defines how quickly the gain reduction can decay back to normal after the trigger that caused the onset of gain reduction has disappeared. Generally speaking, this needs to be more than the minimum setting to avoid distortion. With the attack and release too fast (especially the release), the gain envelope fits around the waveform itself, rather than fitting the general level envelope of the signal, so causing distortion.

RATIO Section

RATIO Fader and Touch Pad

This defines the input to output gain ratio of the compressor, once the signal has reached the level of the threshold value. The Ratio setting is indicated in degrees as an angle, because this allows you to picture the input/output gain characteristic as shown below:

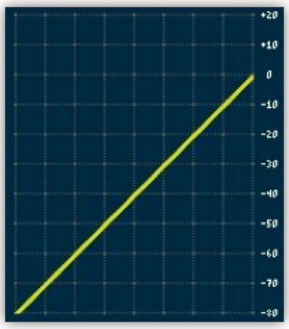

45 degrees
Ratio 1:1
No compression

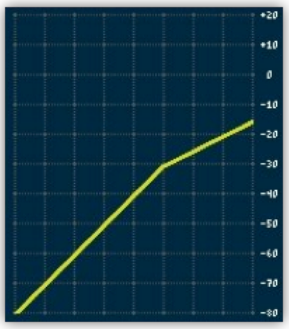

~26 degrees
Ratio 2:1
Gentle compression

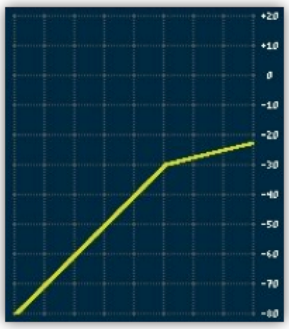

~14 degrees
Ratio 4:1
Significant compression

Note the useful legends to the right of the Ratio fader

For example, when set to 2:1 at around 26 degrees, the compressor gently compresses signals into half the dynamic range they currently occupy above the threshold level.

When set to 0 degrees, the compressor acts as a limiter — the output never goes beyond the threshold level, except for transients that escape the compressor because the attack time is too slow.

The ratio control is not, however, limited to positive values. If you want the compressor to over-react to a breach of the threshold, then set the ratio to a negative value such as -22 degrees; this will roughly double the amount of gain reduction for a given breach of the threshold. For example, if the threshold is 6dB below the peak then, with a ratio of -22, you will get 12dB of gain reduction.

This is useful when you want to eliminate 'esses' when they occur, rather than simply keeping them below a threshold.

SOFT KNEE Button

This control allows for a gradual onset of gain reduction as the signal level approaches the threshold level, rather than a sudden onset of gain reduction when the signal level actually hits the threshold.

The larger the value set here (in 5dB steps up to 20dB), the earlier the onset of gain reduction prior to reaching the threshold. For general frequency specific compression work, this control allows for more natural compression, rather than brick-wall limiting.

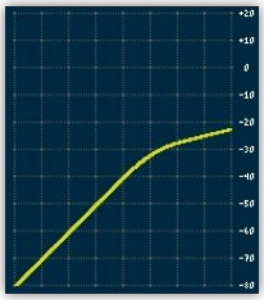

10 dB Soft Ratio

Threshold Section

MAKEUP Touch Pad

This defines how much gain is applied after compression. Makeup is useful in cases where you want to bring the background level up within a specific frequency band, but want the peaks to stay where they are.

4.3 Graphical Screen Controls

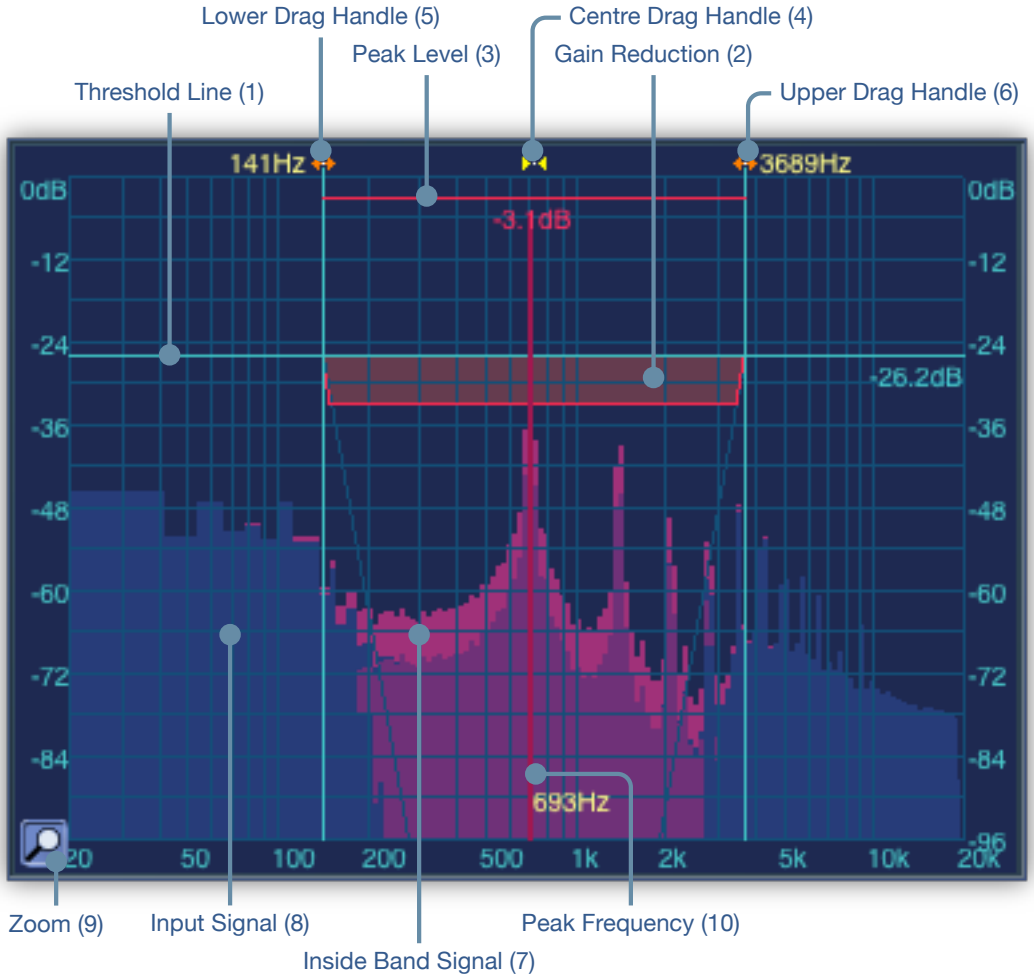

The Oxford SuprEsser's graphical display is the key to its intuitive operation. This display provides all the important pieces of information you need to access visually and quickly, such as the correct threshold and filter frequency values.

In the following descriptions, references are provided to the illustration's annotation in the form **(1)**.

In operation, you will see a FFT (Fast Fourier Transform) of the input signal **(8)** with a pink region above showing the effect of the gain reduction. The FFTs are plotted on a logarithmic scale, and related to actual dB values, so you get a clear idea of how much energy is in a given frequency band.

The FFT shows an important red peak-held vertical line **(10)**. This peak value is useful because it shows you the frequency with the highest energy within the bandpass window. This gives a good indication of where the centre of the 'ess' sound is. Clicking on this peak will centre the window to this frequency, thus

isolating it better.

To help you find the threshold, the most important place to look is the red time-domain peak reading **(3)**. This gives the peak dB level of the audio signal leaving the bandpass filter. Below this, you will see a blue line representing the threshold **(1)**. The distance between these two lines indicates the amount of gain reduction of peaks that you will achieve. With a ratio of 2:1 (the default), if the threshold is 12dB below the peak, you will tend to get 6dB of gain reduction when the peaks occur.

Once you see how it works, you will be able to set the threshold almost without listening!

4.3.1 Explanation of Controls

Band Filter Lower and Upper Cut-off Drag Handles

These can be dragged to move the lower or upper edges of the band filter as desired. Using the Left mouse button, the handles follow the pointer. The Right mouse button, or Shift + Left mouse button, provides a much finer adjustment. The handles sit on top of a vertical blue line that indicates where the lower and upper limits of the band filter are, or to be more precise, the used filters' -3dB cut-off points.

Bandpass Filter Audio Output Peak-Hold Level

The bandpass peak-hold meter shows the peak signal level leaving the bandpass filter, prior to being fed to the Dynamics processing. This is a very important piece of information because the distance between this line and the threshold line (6) is the amount of gain reduction (of peaks) that you will end up with. The hold time is 2 seconds.

Central Filter Drag Handle

This handle allows you to drag the centre frequency of the band window up and down the scale by moving left/right. It also allows you to widen or narrow the band window by dragging up/down. Using Shift + Left mouse button, or Right mouse button, allows you to achieve finer adjustment.

Alternatively, clicking on any arbitrary point on the graph will centre the band filters around this point in the frequency spectrum. Furthermore, if you hold the mouse down, you can then drag the filter up or down, or widen/narrow it in the same way as with the centre drag handle.

This allows you to very slickly widen the window, look at the FFT peaks, and then home in on one (by clicking and dragging down, in one single

movement. This way you can quickly isolate and listen to each of the peaks that occur regularly, to find out, for example, what consonants they represent.

Gain Reduction Indicator

This red-shaded area gives a real-time indication of the shape and amount of occurring gain reduction. The shape follows the slope indicators, and roughly corresponds to the shape of the pink FFT region (see 7 above).

Threshold Line

This line shows the current threshold level. The line will wander up and down as the general signal level rises and falls (unless you have switched the threshold to 'Absolute' mode). If you move the threshold line while there is no signal or a silent signal present, you are setting the presumed level that will occur when the signal returns. This will be indicated on the graph as a dotted line (or ghost threshold line).

Central FFT Peak

This holds for your attention that important frequency in the band window, which has the peak energy. This will most likely indicate a consonant or fricative that you may want to deal with. Clicking on this peak will then centre the band filter to this frequency, and help to isolate it.

Zoom In/Out

This allows you to zoom-in or zoom-out when clicking and dragging. Additionally, you can use the jog/scroll wheel forward/backward to achieve the same thing.

Useful Shortcuts

Jogwheel up/down or forward/backward to zoom in/out. Left/Right-Click anywhere to centre the band filter to that point.

5 Oxford SuprEsser DS (AAX DSP)

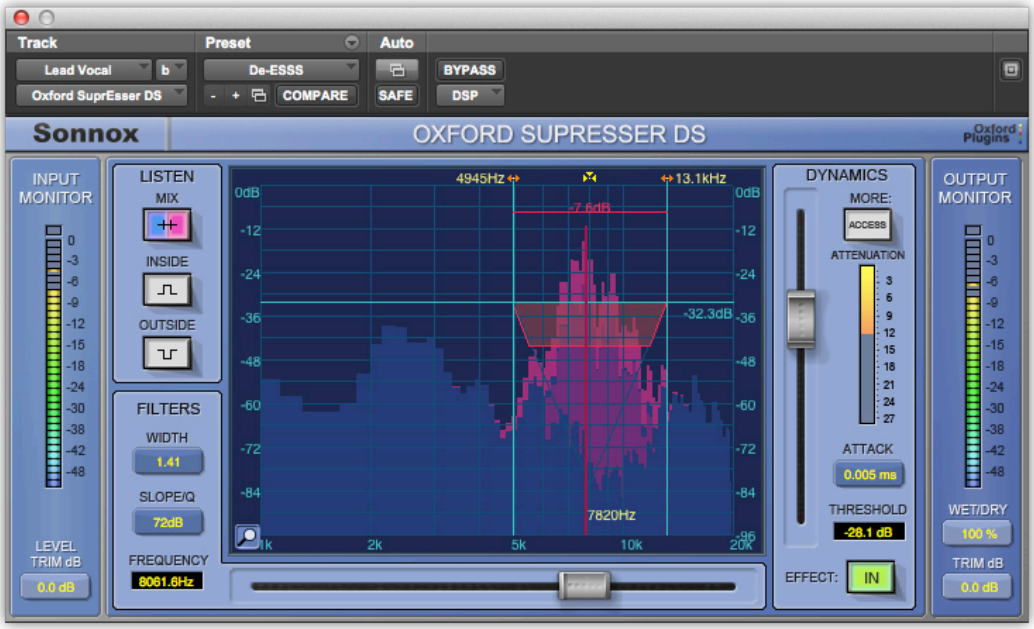

The **Oxford SuprEsser DS** is a highly-featured professional De-Esser AAX DSP plug-in for use with Avid S3L live consoles and Pro Tools|HDX.

In order to provide high-quality processing with low latency, the kernel size has been reduced to 128 samples and the GUI focusses on frequencies greater than 1 kHz.

This enables highly effective deessing with less than 2 ms of delay at 48 kHz.

The SuprEsser DS includes AAX DSP and AAX Native components. This allows offline bouncing in Pro Tools|HDX 11, and session portability between HDX and Native Pro Tools systems (as long as the SuprEsser DS is installed on each system and a **PTXSUPG5** licence is connected).

The DSP and Native components of the SuprEsser DS have different amounts of delay, despite having the same kernel size. Unlike the Native-only SuprEsser (LL, Normal, HR), this delay is constant, and does not change with host buffer size.

DSP delay: 93 samples (1.9 ms at 48 kHz)

Native delay: 180 samples

The Oxford SuprEsser DS is compatible with 44.1 kHz and 48 kHz sample rates.

6 Specifications

Supported Sample Rates	
Native (LL, Normal, HR)	22.05 to 44.1 kHz [1]
Native (LL, Normal, HR)	44.1, 48, 88.2, 96, 176.4, 192 kHz
AAX DSP (SuprEsser DS)	44.1 and 48 kHz
Input-Output	**Parameters**
Input Trim	-12 to 12 dB
Output Trim	0 to -12 dB
Wet/Dry	0 to 100 %
Filters	
Width (LL, Normal, HR)	0.2 to 10 octaves
Width (SuprEsser DS)	0.2 to 4.45 octaves
Slope	12 to 72 dB/octave
Frequency (LL, Normal, HR)	20 Hz to 20 kHz
Frequency (SuprEsser DS)	1 kHz to 20 kHz
Compressor	
Threshold	0 to -84 dBFS
Attack	0.01 to 52 ms
Hold	0.01 to 30 seconds
Release	0.01 to 3.11 seconds
Ratio (degrees)	45 to 0 to -22.5
Ratio	1:1 to Inf:1 to -2.5:1
Gain Makeup	0 to 24 dB
Soft Knee	0, 5, 10, 15, 20 dB
Auto Level Tracking Time Constant	0.1 to 10 seconds
Listen Modes	Mix, Inside, Outside
Processing Modes	
Trigger	Band, Wide
Audio	Band, Wide

[1] At sample rates below 44.1 kHz the internal dynamics algorithm still uses 20 samples of look-ahead. So for samples rates lower than 44.1kHz, you may need to reduce your attack and release times in order for the processing to sound as expected.

6.1 SuprEsser DS – Pro Tools | HDX & S3L – Instances per chip

	48 kHz	96 kHz	192 kHz
Mono	3	–	–
Stereo	2	–	–

7 Preset Manager Toolbar

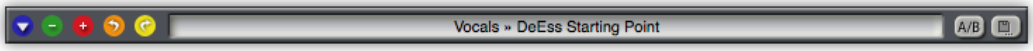

The Oxford SuprEsser plug-in comes equipped with its own onboard Preset Manager, which is displayed as a toolbar at the top of the plug-in window, just as if the host created it (see above). The reasoning behind this is to allow increased portability of your presets across all the host applications, while also providing a consistent and versatile interface. Although most host platforms allow the creation and loading of presets, those host-created preset files are not portable between different host applications. With the Oxford plug-ins' preset manager, you can create a named preset in one host application and load it when using an alternative application. The Preset Manager ToolBar is enabled via the Sonnox button Options Menu (see Description of Controls).

The Sonnox Preset Manager is fully described in a companion document — Sonnox Toolbar and Preset Manager Operation Manual — available for download on the Support Documentation page of www.sonnox.com.

8 Supported Platforms

- Avid Pro Tools (RTAS, AAX Native and DSP 32/64-bit)
- VST hosts (32/64-bit)
- AU hosts (32/64-bit)
- Mac Intel OSX 10.6 or higher
- Windows XP, 7 and 8 (32/64-bit)

9 System Requirements

For latest System requirements, please visit www.sonnox.com.

All versions

- Free iLok account
- Appropriate product licence
- iLok2

Pro Tools

- Approved Digidesign/Avid CPU and hardware configuration
- Pro Tools 8 (Native or HD), or higher

VST Native

- VST compliant host application (e.g. Cubase, Nuendo, etc.)

Audio Units

- Approved Apple CPU and OSX 10.6 or higher
- Audio Unit Host application (e.g. Logic, Digital Performer, etc.)

10 Copyright and Acknowledgements

Trademarks and content copyright © 2007-present Sonnox® Ltd. All rights reserved.

Sonnox® and the five dots logo are registered trademarks of Sonnox Ltd.

This product is manufactured and supplied by Sonnox Ltd. This product is protected by one or more European and/or US patents.

DIGIDESIGN, AVID and PRO TOOLS are trademarks or registered trademarks of Avid Technology Inc.

VST is a trademark of Steinberg AG.

All other product and Company names are trademarks or registered trademarks of their respective holders.